REPORT DOCUMENTATION PAGE

1

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
29-08-2014	Manuscript	-

4. TITLE AND SUBTITLE
Security and Interdependency in a Public Cloud: A Game Theoretic Approach

5a. CONTRACT NUMBER
W911NF-13-1-0157

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
206022

6. AUTHORS
Charles A. Kamhoua, Luke Kwiat, Kevin A. Kwiat, Joon S. Park, Ming Zhao, Manuel Rodriguez

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
Florida International University
11200 SW 8th Street

Miami, FL 33199 -0001

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
62705-CS-REP.6

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT
As cloud computing thrives, many organizations – both large and small – are joining a public cloud to take advantage of its multiple benefits. Especially public cloud based computing, is cost efficient, i.e., a cloud user can reduce spending on technology infrastructure and have easy access to their information without up-front or long-term commitment of resources. Despite those benefits, concern over cyber security is the main reason many large organizations with sensitive information such as the Department of Defense have been reluctant to join a public cloud. This is because different public cloud users share a common platform such as the hypervisor. An attacker

15. SUBJECT TERMS
Cloud computing; cyber security; externalities; game theory; interdependency

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Ming Zhao
UU	UU	UU	UU		19b. TELEPHONE NUMBER 305-348-2034

Report Title

Security and Interdependency in a Public Cloud: A Game Theoretic Approach

ABSTRACT

As cloud computing thrives, many organizations – both large and small – are joining a public cloud to take advantage of its multiple benefits. Especially public cloud based computing, is cost efficient, i.e., a cloud user can reduce spending on technology infrastructure and have easy access to their information without up-front or long-term commitment of resources. Despite those benefits, concern over cyber security is the main reason many large organizations with sensitive information such as the Department of Defense have been reluctant to join a public cloud. This is because different public cloud users share a common platform such as the hypervisor. An attacker can compromise a virtual machine (VM) to launch an attack on the hypervisor which, if compromised, can instantly yield the compromising of all the VMs running on top of that hypervisor. This work shows that there are multiple Nash equilibria of the public cloud security game. However, the players use a Nash equilibrium profile depending on the probability that the hypervisor is compromised given a successful attack on a user and the total expense required to invest in security. Finally, there is no Nash equilibrium in which all the users in a public cloud fully invest in security.

Security and Interdependency in a Public Cloud: A Game Theoretic Approach

Charles A. Kamhoua[1], Luke Kwiat[2], Kevin A. Kwiat[1], Joon S. Park[3], Ming Zhao[4], Manuel Rodriguez[1]
{charles.kamhoua.1; luke.kwiat.ctr; kevin.kwiat; manuel.rodriguez-moreno.1.ctr}@us.af.mil; jspark@syr.edu; ming@cs.fiu.edu
[1]Air Force Research Laboratory, Information Directorate, Cyber Assurance Branch, Rome, NY
[2]University of Florida, Department of Industrial and Systems Engineering, Gainesville, FL
[3]Syracuse University, School of Information Studies (*i*School), Syracuse, NY
[4]Florida International University, School of Computing and Information Sciences, Miami, FL

***Abstract*— As cloud computing thrives, many organizations – both large and small – are joining a public cloud to take advantage of its multiple benefits. Especially public cloud based computing, is cost efficient, i.e., a cloud user can reduce spending on technology infrastructure and have easy access to their information without up-front or long-term commitment of resources. Despite those benefits, concern over cyber security is the main reason many large organizations with sensitive information such as the Department of Defense have been reluctant to join a public cloud. This is because different public cloud users share a common platform such as the hypervisor. An attacker can compromise a virtual machine (VM) to launch an attack on the hypervisor which, if compromised, can instantly yield the compromising of all the VMs running on top of that hypervisor. This work shows that there are multiple Nash equilibria of the public cloud security game. However, the players use a Nash equilibrium profile depending on the probability that the hypervisor is compromised given a successful attack on a user and the total expense required to invest in security. Finally, there is no Nash equilibrium in which all the users in a public cloud fully invest in security.**

Keywords- Cloud computing; cyber security; externalities; game theory; interdependency

I. INTRODUCTION

With software being one of the fastest growing industries in the United States [1], when its security is overlooked the inattentiveness can be attributed to both the producer and consumer. This can have far reaching implications, from infrastructure protection to the home computer system. Internet security suffers too due to under-investment from both sides of the market, which can be counterintuitive since logic dictates that prevailing economic forces should drive the incentive to invest on both ends. This is not the case for several reasons, including perverse incentives, asymmetrical information, and interdependency (we will elaborate on the meanings of these terms from economics in the appropriate parts of our paper.) However, it will be seen that interdependency underpins all these causes and influences network security in general. The preliminary version of this paper was published in [27].

Due to the fast paced nature and rapid expansion of developments in the cyber realm, first mover advantages can be enormous. This can create a software maker philosophy in which "they'll ship it on Tuesday and get it right by version 3" [16]. This philosophy clearly can neglect many security aspects on the supply side. And the demand side, in turn, cannot truly know what it is purchasing, since many of the vulnerabilities could go undetected. This is especially true in large networks with limited security manpower. The idea of "get it out now and fix it later" is a perverse incentive that is created by the demand aspect of the Internet economy since the incentive is to have new and updated versions of software as fast as possible but the unintended is a product that is rife with bugs. However, because of information asymmetry, the consumer usually does not know the true nature of the product he is being delivered. This is because many times the producers do not know the true security of their own product [16]. This is especially true with emerging fields of computing such as cloud computing [15]. And it is indeed a sizable problem, as fears of leakage of sensitive or confidential data poses a "significant barrier to the adoption of cloud services" [17], which hinders major industry players from switching to cloud platform services, stifling its growth. The lack of product knowledge, product testing, and trust all establish an interdependent relationship between producer and consumer. For better or for worse, the feedback mechanism that governs economics is an interdependent relationship between two sides of trade. It allows a consumer to send signals to a producer so that maximum utility can be reached (*i.e.*, Pareto efficiency). However, the examples of perverse incentives and information inequality (where this feedback mechanism has failed in unintended and undesired ways) are just a small part of the general connectedness of network security. In fact, network security is just another small part in the complex infrastructure system of any developed nation. And as we will see, interdependency is the underlying factor in this large network of infrastructures critical to the operations of a country.

The cloud now figures largely in the information infrastructure. It is critical because of its rapidly expanding size and scope. This is especially problematic for the aforementioned problem of technology outpacing security. This spurs cloud providers to furnish expertise in security over what individual organizations (hereby alternately referred to as *users*) would do on their own. This encourages more users to join the cloud; however, the cloud then becomes an attractive target because of the potentially large payoff of a cyber attack. What is more notable than the

Approved for Public Release; Distribution Unlimited: 88ABW-2014-2359, Dated 16 MAY 2014.

regular security issues any network would have is that public clouds exhibit a unique type of interdependency because of the ability of an attacker to propagate his attack through the hypervisor to all VMs using the hypervisor. This eliminates a very important aspect of regular network security in which an attacker would have to go through a multi-hop process in order to launch an indirect attack. Thus, a public cloud at its current stage leaves its users more susceptible to a 'bad neighbor' effect where an unsecured user might allow another to be indirectly attacked. Although our focus is on public clouds, the same research problems may also exist in private clouds, and our solution is also applicable. We focus on public clouds only because the problems are more pronounced in public clouds.

In a dense network of VMs, an attacker may launch an indirect attack on a User *j* by first compromising the VMs of User *i* and then attacking User *j* as a prime target. This creates a risk connection between the users of a cloud where a 'large' player (one who has a high potential loss) may not be willing to use cloud services due to the risk imposed by a 'small' player (low potential loss from a successful compromise). This threat is worsened when a small player will not invest in security measures since it could (correctly) rationalize that an attacker will attack the larger user anyway, so investing would be pointless. Definitely, a single user of a public cloud cannot protect itself if other users are not doing the same. This means that a user will be protected if it defends itself while other users are also securing their asset. When there are two or more rational entities that face interdependent choices, we can use game theory to model their behaviors, as it is indeed "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers" [5].

There are several main contributions this paper makes. Primarily, it aims to model these behaviors that govern the actions of different users in the cloud using game theoretical concepts. Along with modeling the choices of cloud users, it will be shown that the low profile user imposes a negative externality, or a cost imposed unwittingly upon an otherwise uninvolved party—most notably the larger user. This will, in turn, spur the large player to invest more often than the small player since the large player is usually the prime target. The outcome: there is no Nash equilibrium in which all the players will fully invest in security. Lastly, we will prove that the probability that the hypervisor of a cloud is compromised given a successful attack on a VM will determine if we have a pure or mixed strategy Nash equilibrium.

After the related work in Section II, Section III will explain the cloud architecture common to the public cloud model that is incorporated into our game model. Section IV will explain and set up the problem in the context of game theory and diagram the problem in a normal form game. Section V looks at the results provided by the game and maps out the different types of equilibrium reached given different parameters. Further, Section V describes and shows the equilibria changes in accordance with changes to the game parameters. Section VI shows the numerical results that graphically demonstrates how the equilibrium changes following a change in the parameters. Section VII extends the model beyond one attacker and two users so as to pave the way for possible future research in the topic. Section VIII concludes the paper.

II. Background and Related Work

We divide the related work in five subsections. In Subsection *A* we will look at the interdependent nature of the critical infrastructure network in the United States and its connection to cyberspace. In Subsection *B* we will relate game theory and its connection to interdependency. In Subsection *C* we will bring together game theory and network security with no intermediary. With Subsection *D* game theory is applied to cloud computing. Subsection *E* deals with interdependency and cross-side channel attacks between VMs.

A. Critical Infrastructure Defense (and lack thereof)

Generally, the United States government does not interfere in the affairs or operations of the Internet unless it pertains to national security. However, even when national security is at stake, the government is ill-prepared for a response, as Dave Clemente argues in his paper [2]. The main problem, he reasoned in his thesis, is that the infrastructures critical to the operations of the United States are mislabeled and overstated due to miscommunication at the local and national governmental levels. This causes many infrastructures that are not critical to be labeled critical (This is nicely stated in his aphorism: "When everything is critical, nothing is"). The problem is compounded by tying all these infrastructures together through a dense network of interconnectedness, making one network of infrastructure dependent on another. The backbone of this connected network is the Internet, which is becoming increasingly relied upon and only furthering the deep ties these sub-networks already have. Unfortunately, Clemente argues, the Internet securitization process is not keeping pace with the current expansion of the Internet due to industry pressures to sacrifice long term security needs for short and mid-term speed and efficiency needs. And until the critical infrastructure is taken out of private interests (which would cause much more harm than good), this problem will persist. And although no major solution was mentioned by Clemente—other than something must be done— a much more comprehensive solution was laid out by Kenneth Cukier [3].

The work done by Cukier and his colleagues addressed many of the issues raised by Clemente. The main issue was that there is an underinvestment of security within the critical information infrastructure of the United States. This problem was discussed at length and was cast as a symptom, not the disease. The underinvestment was due to many underlying factors such as informational asymmetry (companies do not know the extent of their problem), conflict of interest (government interests vs. private), and interdependent security (this will be further analyzed in the context of game theory later). All these problems aggregate into a general deficiency of investment in cyber security. Although this seems like an economically counterintuitive outcome, it is a

rational one given the constraints of various aforementioned forces. The solution offered by Cukier was essentially an insurance market for security risk, facilitated by a favorable environment created by the government.

Cukier goes on to state that many private companies do not know the extent of their risk because of a reluctance to share their vulnerabilities with others. Insurance companies will not insure the risk since they do not have access to the information to quantify it. This creates a cat-and-mouse game where neither the insurance market nor the companies in need of security will make the first move. This, according to Cukier, is where the government can step in and facilitate transactions of sensitive information as well as preserve anonymity. The creation of a beneficial environment through incentives and information exchange can create a market for risk, which by definition will reduce risk of infrastructure sectors (insurance premiums will discourage risky business and encourage security investing). The dissertation of Forrest Hare [4] reflects these sentiments as he argues that there is an underinvestment due to a conflict of interests. He contends that a public-private partnership should be formed to facilitate the transfer of information and to increase the incentives of private firms to invest in security. This will lead to noticeable positive externalities on the public (since they will be more secure) and everyone will be better off as a result.

Actually, under the new Executive Order 13636—Improving Critical Infrastructure Cyber security [25], the White House would like to provide incentive to private companies to voluntarily adopt a Cyber Security Framework. The Framework is a partnership with the owners and operators of critical infrastructure to improve cyber security information sharing and collaboratively develop and implement risk-based standards. The Framework's goal is to share cyber security information such that the United States government and the private sector may better protect and defend themselves against cyber threats and reduce cyber risk to critical infrastructure. In fact, a security breach on a government contractor (*i.e*. a private company) can compromise multiple government programs. This shows the interdependency between government and private sector security. The White House's Cyber Security Framework is currently under development at the National Institute of Standards and Technology. The Cyber Security Framework includes a set of standards and technological approaches to be adopted by each organization to minimize cyber risks.

B. Game Theory and Interdependency

Through globalization, firms are becoming increasingly dependent upon each other. Thus, it would be logical to assume that their choices would reflect the actions of their competitors and benefactors sharing a given set of information. Game theory accurately describes these conditions, as it is poised "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers" [5]. This makes the case for interdependency among firms, as the actions of one affects the actions of many. The examples of interdependency observed here will include airline security, bankruptcy, and vaccinations.

Two of the papers from the National Bureau for Economic Research (NBER) carefully looked at multiple scenarios involving game theory and the subsequent interdependency of the players [6-7]. The first paper looked at discrete and mostly static games [6]. It was shown that with airline security, one's own investment in baggage security was heavily dependent on the choices of the other airline in a simple two player game. Here one's own security is compromised due to another airline's lack of security or complemented by the reinforcement of the rival's airline security. It was shown that the two Nash equilibria that exist in a simple two firm game occur when both airlines invest in security and when both airlines do not invest in security. As stated in the previous subsection, clearly only the outcome of both investing is desirable. However, economic costs and initial conditions can influence the firms to go the other way and to not investing. With government regulation or other methods to tip incentives toward investing, an economically-optimal situation can be achieved with a little tweaking. Similar results were found with more than two firms since the investing of one firm can cause multiple firms to change their decision to invest, creating a cascade effect in which one firm causes another to invest and so on. Within the same paper [6], similar results were derived from firm bankruptcy. If each division of a large firm, such as bank, were to undergo risk reduction individually, the collective risk of a firm would be reduced. However, if one branch takes exceptional risks, it can create bankruptcy for the whole firm such that the other divisions succumb by the cascading effect.

The second of the NBER papers demonstrated the cascading effect [7]. Again, the airline security problem was studied but in much more depth and mathematical rigor. They proved that the incentive to invest is heavily dependent on the cost of investing compared to the benefit derived from both investing in security. The cost could be manipulated both by lowering the cost of investing as well as raising the cost of not investing.

Unlike an organization having exclusive use of computational resources, the resource sharing that occurs in the cloud enables unforeseen exploitation of weaknesses by attackers. Similarly, the commonality of computational resources without an equal commonality of user-instantiated security creates an avenue for launching an attack on other tenants i.e., a negative externality due to interdependency and resource sharing.

The link between interdependency and game theory has been clearly established along with the connection between network security and interdependency in the previous two subsections. In the next section we will show the application of game theoretical concepts to network security.

C. Applying Game Theory to Cyber Security

Sun *et al*., presented a model of investment security [8] where they simulated a security game between two arbitrary companies having to decide whether to invest or not invest in information security. The payoffs were based on several

inputs such as cost of investing and the possible loss from a security compromise. However, the most important parameter discussed was a penalty parameter *p* for not investing. It was shown that the 3 Nash equilibrium strategies produced from the game were two pure Nash equilibria (neutral payoff for not investing and a positive payoff each for investing) and one mixed strategy that was a function of all the parameters. The pure strategies were shown to have an Evolutionary Stable Strategy (ESS) while the mixed strategy was not. The mixed strategy was demonstrated to be a focal point, as a strategy on either side of this critical point 'tipped' or 'cascaded' to the closer ESS at pure Nash equilibrium. However, *p* was shown to factor in where the mixed strategy fell between the two pure strategies on the probability spectrum of 0 to 1. This could skew the results from what could be considered 'normal' and demonstrated that an outside force such as the government could manipulate the penalty parameter in order to achieve a more favorable outcome.

Even though the previous example would have used a central manager or network administrator to decide if investing was the correct choice, Kamhoua *et al*. applied game theory to nodes in autonomous networks [9]. They used similar constraints with similar results: there are 3 Nash equilibria, two pure and one mixed with the mixed strategy being an unstable equilibrium. This resulted in a cascading of strategies of either side to that tended toward the two pure Nash equilibria. The main difference, however, is instead of a penalty parameter, as in the paper of Sun *et al* [8], there is a trust parameter which the initial conditions of the strategy heavily depended on. The trust parameter depended on how much the deciding node believes that other node will participate in a security mechanism. The main conclusion to draw from these simulations is that it is impossible to move from the low trust equilibrium to the high trust equilibrium through an evolutionary process. In the replicator dynamic model [26], the final state depends entirely on the initial condition. This has broad reaching implications, from network security to cloud computing.

In Tamer Basar's and Tansu Alpcan's book [10], they explain the devastating costs of failure to properly protect a network. They show how an attacker can infiltrate a network at one node, but spread to other nodes (or infrastructures) due to contagion. This can cause a spillover effect where one node affects another and so on. The end result is that network interdependency is created and that one unprotected node causes risks at all the other nodes, so the decision of one affects the outcomes of many. Basar and Tansu however only applied network security in a traditional computer setting. The rise and expansion of cloud computing has led to many questions about its security. To raise concerns further, cloud computing's annual growth is rapidly outpacing regular computing methods by a significant margin [11]. In the next subsection we will outline details on its expansion, tradeoffs in switching to cloud platforms, and further research in cloud security.

D. Interdependency Analysis in Cloud Computing

According to the National Institute of Standards and Technology definition of cloud computing, some of the 'essential characteristics' that come with the term include resource pooling, elasticity, resource optimization, network access and on-demand self-service [12]. Though this can overcome many constraints posed by traditional computing, the emerging field of cloud computing currently carries some profound tradeoffs. Pearson and Benameur outlined several important drawbacks in cloud technology such as privacy, security, and trust concerns [13]. However, these three problems are not unrelated to each other. Security within the cloud is based on trust of the provider, and privacy is based on the relevant security issues. Trust is in turn built on the relationship of security and privacy that the cloud operator provides. This is not the case every time, since not all cloud technology has these aforementioned problems due to their diverse nature. Zissis *et al*. [14] differentiate between public and private cloud structures by stating that private cloud technology is for inter-organizational operations and no third party is required while public and community cloud computing utilizes a third party for a variety of service platforms. Such service platforms that cloud computing provide include Infrastructure as a Service (IaaS), Software as a Service (SaaS), and Platform as a Service (PaaS).

An IaaS cloud provides a user access to virtualized hardware, presented by a hypervisor (*e.g*., VMware, Xen, KVM) and encapsulated in a VM, where the user is able to deploy and run arbitrary software including operating systems and applications on the underlying shared hardware. A PaaS cloud provides a user a language-specific platform (*e.g*., JVM, .Net) to deploy and run arbitrary applications developed using the given language on the underlying shared platform. A SaaS cloud provides a user access to a particular application (*e.g*., web-based email, document editor) where the user can use the functionality provided by the underlying shared application. Although these different levels of cloud services can be built separately, it is increasingly common to build a high-level cloud service using resources provided by a lower-level one (*e.g*., build a SaaS on resources from PaaS and a PaaS on resources from IaaS), so that the former can benefit from the elasticity and economics provided by the latter. Therefore, although our paper focuses on VM-based hosting of mission-critical applications in an IaaS setting, its outcomes can also generate an impact to other models of cloud computing (further information can be seen in [14]). Although private clouds do share some of the benefits and drawbacks of public clouds, the issues of privacy, security, and trust arise from mainly public cloud platforms, as many of the users' computing capabilities are outsourced to a third party owner who leases the technology in a variety of ways. Therefore we focus on the public cloud; so in this paper private cloud entities will not be discussed further. In fact, private clouds allow users from the same organization to run their internal applications on shared resources. Therefore, in a game theoretic sense, there should be less conflict of interest among private cloud users since they belong to the same organization.

As stated before, these problems that involve the public cloud are not unrelated as they all underpin a unique relationship between the third party provider and the cloud user. This can give rise to interdependency between the user and the operator of the cloud. If we apply the behavior of network nodes as described in [9] to a cloud's VMs, then we can see that cloud computing yields very interdependent structure. Cloud computing gives way to two types of interdependent relationships: cloud host-to-client and cloud client-to-client.

Client-to-client interdependency is much less studied than to the above-mentioned cloud host-to-client relationship. Although, it can still carry the negative externalities provided by the first relationship since a security compromise is the same no matter where it has originated. A simple example of this involves the airline security problem found in [6] and [7] where a bomb infused baggage is sent through an unsecured airline, which in turn reaches a heavily secure airline because no inter-airline security screening is used (and it usually is not). Thus, an under-secure airline can impose negative externalities onto a seemingly secure airline. Similarities can be drawn to two clients operating in the same cloud environment. An attacker can compromise an unsecured client and make its way to the more secure and larger client through the hypervisor. However, unlike the airline interdependent security problem where a bomb can only destroy one airline, a virus in a public cloud or computer network can compromise many VMs including the VM in which the attack originated.

We have already seen that interdependency lays the foundation for game theory in previous subsections. Indeed, this scenario between two clients also involves two or more intelligent rational entities with conflicting incentives. Analogous to the previous example, a small firm with high overhead will see little point to invest in security since its cost to invest is most likely diminished by the fact it has lower possible loss from being compromised. However, a larger firm has a much higher potential loss from being compromised, especially if they carry sensitive information (This has been seen in [15] when large firms refuse to use cloud computing because of its risks). Thus, a rational attacker might attack a smaller firm, compromise the hypervisor, and then target the larger firm if the potential gain from a successful indirect attack outweighed the potential gain from a direct attack.

E. Interdependency and Cross-side Channel Attacks between VMs

The support for security isolations from existing cloud systems is limited. The different VMs sharing the same resources may belong to competing organizations as well as unknown attackers. From the perspective of a cloud user, there is no guarantee whether the underlying hypervisor or the co-resident VMs are trustworthy. The shared resource makes privacy and perfect isolation implausible. There is a risk that a covert side channel be used to extract another user's secret information or launch a Denial of Service (DoS) attack. Cross-side channel attacks between VMs are possible in a public cloud when the VMs share the same hypervisor, CPU, memory, and storage and network devices. Some of the resources can be partitioned (*e.g.*, CPU cycles, memory capacity, and I/O bandwidth). VMs also share resources that cannot be well partitioned such as last-level cache (LLC), memory bandwidth, and IO buffers. The shared resources can be exploited by attackers to launch cross-side channel attack. Although a multi-tenant public cloud-computing environment provides various advantages, it also introduces new challenges and concerns, especially on security issues. For instance, the security problems on a shared cloud resource (e.g., cloud storage devices, network services, software components, etc.), which are originally rooted from one of the tenants via internal vulnerabilities or external cyber-attacks, may eventually affect the service quality and security of all the tenants in the same cloud-computing environment. Unfortunately, we cannot simply assume that there would be a single authority who could comprehensively maintain all the possible issues, not only technical but also non-technical, across the tenants.

Moreover, existing cloud service providers do not provide sufficient security guarantees to their tenants. In fact, the service-level agreements (SLAs) of representative cloud providers (e.g., Amazon EC2/S3, Windows Azure, Google Compute Engine) specify only the provisions related to service up time, and there is no mentioning of security in these SLAs at all.

Many researchers have investigated the cache based side channel. Ristenpart *et al*. [18] show that a malicious user can analyze the cache to detect a co-resident VM's keystroke activities and map the internal cloud infrastructure and then launch a side-channel attack on a co-resident VM. Bates *et al*. [19] demonstrate the ability to initiate a covert channel of 4 bits per second, and confirm co-residency with a target VM instance in less than 10 seconds. Li *et al*. [23] proposed several techniques to protect VMs from untrusted management VM, which includes modifying the hypervisor to restrict access of the privileged domain to the memory mappings of the VM, encrypting all of the memory pages and vCPU registers before they are accessed by the privileged domain, and providing a hash value of the kernel image to be compared with the one residing on the VM. HyperSentry [24] enables stealthy in-context measurement of hypervisor integrity using a hardware channel to trigger the measurement and, using the system management mode, to protect the measurement agent's base code and critical data.

Given the danger of a cross-side channel attacks, some users may require physically isolated resources from the cloud provider. Zhan *et al*. [20] introduce HomeAlone - a defensive tool that helps users determine if their VMs have an exclusive use of a physical machine. HomeAlone can detect the activity of an intruder's co-resident VM by analyzing a portion of the L2 memory cache set aside by his VMs. The same technique can be used to detect adversarial VMs which try to extract information through the side channel due to their usual cache activity pattern. This solution, however, requires that all the user VMs to be co-resident which is often difficult to achieve and makes them more vulnerable to hardware and hypervisor failures.

Approaches that dedicate a physical machine to a specific user also greatly limit some of the benefit of a public cloud such as the on-demand dynamic resource allocation. This means that a user can no longer purchase exactly the capacity they require when they require it. Therefore, we consider in this paper only schemes in which the VMs from different users share the same resources. We can see that a cross-side channel attack between VMs is closely related to the problem of interdependency when many users share the same resource that they depend on. This paper provides a comprehensive analysis of direct vs. indirect attack, collateral damage, and negative eternality in a public cloud.

III. System Model

Figure 1 illustrates our system model: A public cloud with n users that we denote User 1, User 2 … User n. Each user runs several applications illustrated by Application 1 …Application k in Fig. 1. Technically, the users may run a different number of applications without any impact on this model. The different applications require an operating system to function and that operating system, in turn, manages a VM in the cloud. In practice, a single user may use several operating systems or numerous VMs.

However, we consider the architecture in Fig. 1 to simplify the exposition. As it is a common practice in a public cloud, we consider that the different VMs from the different users share the same hypervisor and hardware as in Fig. 1. The hypervisor can be of different types such as the Kernel-based Virtual Machine (KVM), Xen, and VMware. The common factor is that the VMs share the same platforms that expose each user to collateral damage.

We consider the possibility of a random hardware failure to be a rare event and neglect that possibility in our analysis. It is well known that the users' security heavily depends on the cloud provider. We are analyzing security interdependency among the users; therefore our model considers that the attacker compromises the hypervisor in two steps. The first step is to compromise a user's VM, or masquerade as legitimate user to obtain a VM in the public cloud. The second step is to use the compromised VM to attack the hypervisor. This means that the public cloud provider takes all the necessary measures to prevent an attacker from directly compromising the hypervisor without using a compromised VM. This is to separate cloud client-to-client interdependency and cloud host-to-client interdependency. However, any model that analyzes cloud host-to-client interdependency can be superposed to our model. We distinguish two types of attack depending on the extent of the consequence: a restricted attack and an unrestricted attack. A restricted attack on User i only compromises the applications, operating system and VM that belong to User i; the hypervisor is not affected after a restricted attack. An unrestricted attack has consequences that can cross a VM to reach the hypervisor, *i.e.* the hypervisor is compromised. We consider that all the users suffer the consequences (damage) if the hypervisor is compromised. This is because an attacker that compromises the hypervisor can then compromise all the VMs on that public cloud.

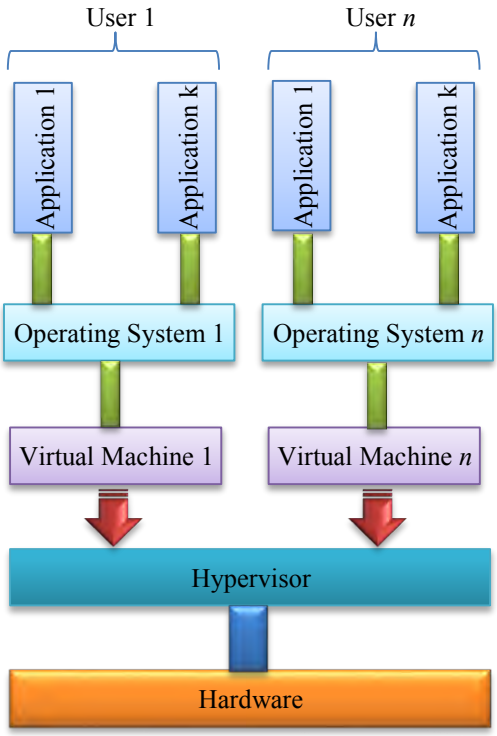

Figure 1: System Model Illustration

We can see that an unrestricted attack causes collateral damage. A direct attack on User i can go through that user's VMs to compromise the hypervisor and ultimately affect the VM of another User j. We also refer to this as an indirect attack on j. Thus, each user in a public cloud can suffer from two types of attack. A direct attack on a User i is when an attacker primary target is User i. Furthermore, an indirect attack on User i happens when an attack that is launched on another User j compromises the hypervisor before compromising User i's VM.

This system model clearly shows that cyber security in a public cloud depends not only on a particular user but also on any other user of the cloud. This is the problem of interdependency. Section IV will analyze the interdependency problem from a game theoretic perspective.

IV. Game Model

This section considers a game with three players: An attacker and two users (User i and User j). Section VII will extend this model to more than two users and multiple attackers. The three players are assumed to be rational, which means that each player has an understanding of the system and has the ability to perform the necessary calculation to only take the actions that maximize his expected payoff. The attacker has two strategies: launch an attack on User i (A_i) and launch an attack on User j (A_j). The attacker can only use one of the two strategies at a time. The attacker strategy to launch an attack on User i may consist of a multi-stage process involving steps such as scanning, collecting information, credential compromising, executing attack payload, establishing backdoor, cleaning footholds

and avoiding firewalls. Choosing to invest is a binary decision for each user in which the two users can either *Invest* (*I*) in security to maintain a minimum security standard and increase their protection or *Not invest* (*N*), *i.e.,* there is no partial investment in security. The strategy Invest may consist of multiple actions such as system monitoring, reconfiguration, patching, updating software, and buying a new antivirus. Investment in security requires a total expense *e*. A strategy profile is a 3-tuple that indicates the action of each player. For instance, the strategy profile (N, I, A_j) indicates that User *i* does not invest (*N*), User *j* invests (*I*), and the attacker launches an attack on User *j* (A_j).

The probability of a successful attack on a user, given that he has invested in security, is q_I and the probability of a successful attack on a user, given that he has not invested, in security is q_N. We assume that

$$0 \leq q_I < q_N \leq 1. \quad (1)$$

We have $q_I < q_N$ because any rational user will only invest in security measures that diminish his chance to get compromised.

The probability that the hypervisor is compromised given a successful attack on a user is denoted π . Our model considers that at least some successful attack on a VM will reach the hypervisor or that $\pi > 0$. In fact $\pi = 0$ means that a successful attack on a VM would never reach the hypervisor which would be a strong assumption. We also consider that not all the successful attacks on a VM can compromise the hypervisor ($\pi < 1$). Thus we have

$$0 < \pi < 1. \quad (2)$$

We consider that there is a high profile User *j* and a low profile User *i*. In case of a security breach, the high profile user incurs more loss than the low profile user. The high profile User *j's* expected loss from a security breach is L_j and the expected loss from User *i* is L_i. Then we consider that

$$0 < L_i < L_j. \quad (3)$$

We will show that this imbalance affects the investment decision of each player and may yield positive and negative externalities. A positive (negative) externality is an action of a player that transfers a positive (negative) effect onto a third party. In fact, when (high profile) users in a public cloud invest in security to protect their applications, operating systems and VMs, they also protect the hypervisor which in turn protects other users from an indirect attack or cross-side channel attack. This yields a positive externality to other users in a public cloud. On the contrary, if a (low profile) user chooses not to invest in security, then an easy attack path to the hypervisor is created and thus exposes all other users of a public cloud to a cross-side channel attack. This yields a negative externality to other users in a public cloud.

The accuracy of our model depends on the correct estimation of the probabilities q_I, q_N, π and the loss $L_i$and L_j. We propose two different approaches to estimation. The first approach is the QuERIES approach [21]. The QuERIES approach estimates the probabilities and costs of successful attacks by first building an attack graph represented as a Partially Observable Markov Decision Process (POMDP). Then QuERIES uses a controlled red-team experiment and information market mechanisms to estimate the POMDP parameters. The outcome of an information market is a collective estimate of a quantity. The red-teams have real financial incentives for making correct predictions of the POMDP probabilities. Finally, the POMDP's optimum policy is calculated to derive the different probabilities and cost.

The second approach to estimate the relevant probabilities and cost associated with our model is based on historical data. In fact, In October 2011, the United States Securities and Exchange Commission (SEC) issued a new guidance [22] requiring that companies disclose cyber incidents including a description of the costs, other consequences, and the relevant insurance coverage. Those data can now be aggregated to estimate the relevant probabilities and cost associated with our model.

In addition, each user has a reward *R* from using the cloud computing services. The reward *R* can be calculated as a function of a user's multiple benefits of using the cloud such as: reduced spending on technology infrastructure; easy access to their information without up-front or long-term commitment of resources; and dynamically grow and shrink the resources provisioned to an application on demand.

Finally, we consider that a User *i* can detect and identify a co-resident VM from User *j* in the cloud via side-channel analysis as in HomeAlone [20]. Further, a skillful attacker will first scan a public cloud to learn about the different users – gaining knowledge of their weaknesses and vulnerabilities before launching an attack. Also, each of the following can be made known or can be estimated about a player [21-22]: the expected loss from a security breach and the related probability; the total expense required to invest in security; and the reward from using the cloud. Therefore, our model assumes that the player's identity, strategy and payoff are common knowledge among the players.

TABLE I: GAME MODEL IN NORMAL FORM

		Attack *i*	
		User *j*	
		I	*N*
User *i*	*I*	$\{R - e - q_I L_i;$ $R - e - q_I \pi L_j;$ $q_I L_i + q_I \pi L_j\}$	$\{R - e - q_I L_i;$ $R - q_I \pi L_j;$ $q_I L_i + q_I \pi L_j\}$
	N	$\{R - q_N L_i;$ $R - e - q_N \pi L_j;$ $q_N L_i + q_N \pi L_j\}$	$\{R - q_N L_i;$ $R - q_N \pi L_j;$ $q_N L_i + q_N \pi L_j\}$

		Attack *j*	
		User *j*	
		I	*N*
User *i*	*I*	$\{R - e - q_I \pi L_i;$ $R - e - q_I L_j;$ $q_I \pi L_i + q_I L_j\}$	$\{R - e - q_N \pi L_i;$ $R - q_N L_j;$ $q_N \pi L_i + q_N L_j\}$
	N	$\{R - q_I \pi L_i;$ $R - e - q_I L_j;$ $q_I \pi L_i + q_I L_j\}$	$\{R - q_N \pi L_i;$ $R - q_N L_j;$ $q_N \pi L_i + q_N L_j\}$

Table I shows the game model in normal form. We can see that Table I is a combination of two tables (left and right). The left table shows the game model when the attacker targets User *i*. Therefore, User *j* can only be subject to collateral damage after a successful attack on User *i* and compromising of the hypervisor (which can happen with probability $q_I\pi$ if User *i* invests or probability $q_N\pi$ if User *i* does not invest). Similarly, the right table shows the game model when the attacker targets User *j* and User *i* can only be subject to collateral damage. The fourth line in each table shows when User *i* chooses to invest while the fifth line shows when User *i* chooses not to invest. In each table, the decision of User *j* is represented in the third (Invest) and fourth (Not invest) column. The payoffs in each block are represented in three lines. The first line is User *i's* payoff. The second line is User *j's* payoff. The attacker payoff is represented in the third line.

The payoffs are calculated as follows: If the player chooses the strategy profile (*I, I,* A_i), both users invest (play *I*) while the attacker targets User *i* (A_i) (left table, fourth line, third column). Then both users get the reward *R*. Both users incur expense *e* because both of them have invested in security. Since the attacker targets User *i* that will be compromised with probability q_I (because User *i* has invested), it will incur a loss L_i if compromised. This will result in an expected loss of $q_I L_i$. User *j* is not targeted but can incur a loss L_j if the attack on User *i* is successful (which happens with probability q_I) and the hypervisor is compromised (which happens with probability π). This is an expected loss of $q_I\pi L_j$ and can also be called collateral damage or loss from an indirect attack. The attacker's payoff is the sum of the expected loss of all the users$U_a(I,I,A_i) = q_I L_i + q_I\pi L_j$. The attacker's partial payoff $q_I L_i$ comes from a direct attack on User *i* while the second part of his payoff $q_I\pi L_j$ is the result of an indirect attack on User *j* through the hypervisor.

However, in the strategy profile (*N, I,* A_i), User *i* has not invested (*N*), User *j* has invested (*I*) and the attacker targets User *i* (A_i) (left table, fifth line, third column). The User *i* does not incur any expense *e* because the user has not invested in security. However, his likelihood of being compromised increases to q_N. Moreover, although User *j* has invested in security, his potential losses from collateral damage increase to $q_N\pi L_j$. The difference $q_I\pi L_j - q_N\pi L_j = (q_I - q_N)\pi L_j$ is a negative externality that User *i* imposes on User *j* by not investing while User *i* is the prime target of the attacker. The attacker's payoff is $U_a(N,I,A_i) = q_N L_i + q_N\pi L_j > q_I L_i + q_I\pi L_j = U_a(I,I,A_i)$. The inequality holds because of (1). The players' payoffs in the other six strategy profiles are calculated in a similar way.

V. GAME ANALYSIS

The main goal of this analysis is to derive the different Nash equilibria of the game in Table I and understand their consequence for both users. At a Nash equilibrium profile, no player's payoff can be increased by a unilateral deviation. Also, each player is playing his best response to other players' best strategies. Therefore, the Nash equilibrium can help predict the behavior of any rational player (*i.e.*, that want to maximize their payoff in a game).

We observe that a user that is the prime target must be hurt before the other user suffers any collateral damage. Recall that the prime target's VM must be compromised before the hypervisor is compromised. Thus, we consider in the remainder of this analysis that each user prefers to invest instead of not investing when he believes that he is the attacker's prime target. For User *i* this translates to

$$R - e - q_I L_i \geq R - q_N L_i \Rightarrow$$
$$e \leq (q_N - q_I)L_i \qquad (4)$$

Similarly, for User *j* this translates to

$$R - e - q_I L_j \geq R - q_N L_j \Rightarrow$$
$$e \leq (q_N - q_I)L_j \qquad (5)$$

Also observe that investing in security is the best option to either User *i* or User *j* if and only if the user believes that he will be the attacker's prime target. Also, the attacker targets only the player that gets him the higher total payoff (consisting of a direct and indirect payoff).

Theorem 1:

If $\pi \leq \pi_0 = \frac{q_I L_j - q_N L_i}{q_N L_j - q_I L_i}$, then the game in Table I admits a pure strategy Nash equilibrium profile (N, I, A_j).

If $\pi > \pi_0$, there are three possible mixed strategy Nash equilibria depending on the required expense for security *e*.

Proof:

We start to analyze the eight different pure strategy profiles to see if one can be a Nash equilibrium.

Case 1: Both users invest,

$$U_a(I,I,A_j) - U_a(I,I,A_i) =$$
$$(q_I\pi L_i + q_I L_j) - (q_I L_i + q_I\pi L_j) = q_I(1-\pi)(L_j - L_i).$$

Then by considering (2) and (3) we have

$$U_a(I,I,A_j) - U_a(I,I,A_i) = q_I(1-\pi)(L_j - L_i) > 0. \qquad (6)$$

Therefore, the attacker gets a higher payoff by targeting User *j* when both users invest. Thus the strategy profile (I,I,A_i) can never be a Nash equilibrium because the attacker can increase his payoff by changing his strategy to A_j. This gets us to the strategy profile (I,I,A_j) that cannot also be a Nash equilibrium because User *i* (not being the attacker's prime target) can increase his payoff by changing his strategy from *I* to *N*. This yields the strategy profile (N,I,A_j) that we study in Case 4 below.

Case 2: Both users do not invest,

$$U_a(N,N,A_j) - U_a(N,N,A_i) =$$
$$(q_N\pi L_i + q_N L_j) - (q_N L_i + q_N\pi L_j) = q_N(1-\pi)(L_j - L_i).$$

Then by considering (2) and (3) we have

$$U_a(N,N,A_j) - U_a(N,N,A_i)$$
$$= q_N(1-\pi)(L_j - L_i) > 0. \qquad (7)$$

Thus, the attacker gets a higher payoff by targeting User *j*. The strategy profile (N,N,A_i) cannot be Nash equilibrium because the attacker can increase his payoff by changing his strategy to A_j. This gets us to the strategy profile (N,N,A_j) that cannot also be a Nash equilibrium because User *j* being the attacker's prime target can increase his payoff by changing his strategy from *N* to *I* (because of (5)). This

yields again the strategy profile (N, I, A_j) that we study in Case 4 below.

Case 3: User *i* invests while User *j* does not.

We can see from Table I that

$$U_a(I, N, A_i) = U_a(I, I, A_i) = q_I L_i + q_I \pi L_j. \quad (8)$$

Moreover,

$$U_a(I, N, A_j) - U_a(I, I, A_j) = (q_N \pi L_i + q_N L_j) - (q_I \pi L_i + q_I L_j) \Rightarrow$$
$$U_a(I, N, A_j) - U_a(I, I, A_j) = q_N(L_j + \pi L_i) - q_I(L_j + \pi L_i) = (q_N - q_I)(L_j + \pi L_i) > 0. \quad (9)$$

Note that the last inequality in (9) holds because of (1).

Combining (8) and (9) we have

$$U_a(I, N, A_i) = U_a(I, I, A_i)$$

and

$$U_a(I, N, A_j) > U_a(I, I, A_j) \Rightarrow$$
$$U_a(I, N, A_j) - U_a(I, N, A_i) > U_a(I, I, A_j) - U_a(I, I, A_i)$$

Taking (6) into consideration we have

$$U_a(I, N, A_j) - U_a(I, N, A_i) > 0. \quad (10)$$

From (10), the attacker gets a higher payoff by targeting User *j*. Thus the strategy profile (I, N, A_i) cannot be Nash equilibrium because the attacker can increase his payoff by changing his strategy to A_j. This gets us to the strategy profile (I, N, A_j) that also cannot be a Nash equilibrium because User *j* (being the attacker's prime target) can increase his payoff by changing his strategy from *N* to *I* (because of (5)). We come back to the strategy profile (I, I, A_j) that we study in Case 1 above, which finally yields Case 4 below.

Case 4: User *j* invests while User *i* does not.

$$U_a(N, I, A_j) - U_a(N, I, A_i) = (q_I \pi L_i + q_I L_j) - (q_N L_i + q_N \pi L_j)$$
$$= (q_I L_i - q_N L_j)\pi + (q_I L_j - q_N L_i) = f(\pi)$$

We can see that $f(\pi)$ is a linear function with slope $(q_I L_i - q_N L_j)$ and initial value $(q_I L_j - q_N L_i)$. From (1) and (3) we have the slope $q_I L_i - q_N L_j < 0$. Thus, $f(\pi)$ is decreasing. Moreover, there is a unique value of π such that

$$f(\pi) = 0 \Rightarrow \pi = \pi_0 = \frac{q_I L_j - q_N L_i}{q_N L_j - q_I L_i}, \quad (11)$$

Furthermore, we have $f(\pi) > 0$ for $\pi < \pi_0$ and $f(\pi) < 0$ for $\pi > \pi_0$. Also,

$$f(1) = (q_I L_i - q_N L_j) + (q_I L_j - q_N L_i) = (q_I - q_N)(L_i + L_j) < 0. \quad (12)$$

The last inequality holds because of (1).

In addition, the initial value is

$$f(0) = q_I L_j - q_N L_i, \quad (13)$$

which can be either negative or positive. Observe that because of (2) the condition $\pi \le \pi_0$ can hold if $0 < \pi_0 < 1$, and by the Intermediate Value Theorem, and based on (12) and (13), it is only possible when $f(0) > 0 \Rightarrow q_N L_i < q_I L_j \Rightarrow$

$$L_i < \frac{q_I}{q_N} L_j. \quad (14)$$

Then we can distinguish two subcases (4a) and (4b).

Subcase (4a): If $\pi \le \pi_0$, then we have $U_a(N, I, A_j) - U_a(N, I, A_i) \ge 0$. Thus the attacker prefers to attack User *j* than to attack User *i*. User *j* prefers to invest than not to invest (see (5)). User *i* not being the attacker's prime target prefers not to invest. Then the strategy profile (N, I, A_j) is the pure strategy Nash equilibrium of the game because no player can increase his payoff by a unilateral deviation.

Subcase (4b): If $\pi_0 < \pi$ (regardless of the sign of $f(0)$) we have $f(\pi) < 0$ and then $U_a(N, I, A_j) - U_a(N, I, A_i) < 0$. The attacker prefers to attack User *i* than to attack User *j*. Thus the strategy profile (N, I, A_j) cannot be Nash equilibrium because the attacker can increase his payoff by changing his strategy to A_i. This gets us to the strategy profile (N, I, A_i) that also cannot be a Nash equilibrium because User *i* being the attacker's prime target can increase his payoff by changing his strategy from *N* to *I* (see (4)). This brings us to the Case 1 above which you recall brings us to Case 4. Therefore, this circular reasoning tells us that there is no pure strategy Nash equilibrium. However, there will be a mixed strategy Nash equilibrium that we analyze next.

Mixed Strategy Nash Equilibrium:

To find the mixed strategy Nash equilibrium, we set three variables α, β, λ with

$$0 \le \alpha, \beta, \lambda \le 1. \quad (15)$$

α represents the probability by which the User *i* plays *I*. Since User *i* has only two strategies, User *i* plays *N* with probability $1 - \alpha$. Similarly, User *j* plays *I* with probability β and plays *N* with probability $1 - \beta$. Likewise the attacker attacks *j* with probability λ and attacks *i* with probability $1 - \lambda$.

By definition, User *i* plays a mixed strategy if and only if his payoff $U_i(I)$ when playing *I* is equal to his payoff $U_i(N)$ when Playing *N*. This translates to:

$$U_i(I) = U_i(N) \Rightarrow (1-\lambda)\beta(R - e - q_I L_i) + (1-\lambda)(1-\beta)(R - e - q_I L_i)$$
$$+\lambda\beta(R - e - q_I \pi L_i) + \lambda(1-\beta)(R - e - q_N \pi L_i) =$$
$$(1-\lambda)\beta(R - q_N L_i) + (1-\lambda)(1-\beta)(R - q_N L_i)$$
$$+\lambda\beta(R - q_I \pi L_i) + \lambda(1-\beta)(R - q_N \pi L_i)$$
$$\Rightarrow \lambda = \lambda_i = \frac{(q_N - q_I)L_i - e}{(q_N - q_I)L_i}. \quad (16)$$

Equation (4) shows that $0 \le \lambda_i \le 1$. Also,

$$U_i(I) < U_i(N) \Rightarrow 0 \le \lambda_i < \lambda \le 1, \quad (17)$$

and

$$U_i(I) > U_i(N) \Rightarrow 0 \le \lambda < \lambda_i \le 1. \quad (18)$$

This means that, if the attacks on User *j* are more frequent than λ_i(and then User *i* is attacked less often), then User *i* prefers to play *N*. User *i* plays *I* otherwise.

Similarly, User *j* plays a mixed strategy if and only if his payoff $U_j(I)$ when playing *I* is equal to his payoff $U_j(N)$ when playing *N*. This translates to:

$$U_j(I) = U_j(N) \Rightarrow \lambda = \lambda_j = \frac{e}{(q_N - q_I)L_j}. \quad (19)$$

Equation (5) shows that $0 \le \lambda_j \le 1$. Also,

$$U_j(I) < U_j(N) \Rightarrow 0 \le \lambda < \lambda_j \le 1, \quad (20)$$

and

$$U_j(I) > U_j(N) \Rightarrow 0 \le \lambda_j < \lambda \le 1. \quad (21)$$

Further, the attacker plays a mixed strategy if and only if his payoff $U_a(A_i)$ when attacking User i is equal to his payoff $U_a(A_j)$ when attacking User j. This translates to:

$$U_a(A_i) = U_a(A_j) \Rightarrow \beta(L_j + \pi L_i) - \alpha(L_i + \pi L_j) = \left(\frac{q_N}{q_N - q_I}\right)\left[(L_j + \pi L_i) - (L_i + \pi L_j)\right]. \quad (22)$$

Given the condition in (16), (19) and (22), we can distinguish three cases that we denote M1, M2 and M3 depending on if $\lambda_j = \lambda_i$, $\lambda_j < \lambda_i$, or $\lambda_j > \lambda_i$. Furthermore, we will see that the total expense required to invest in security e determines which of the mixed strategy is used.

Case M1: If $\lambda_j = \lambda_i \Rightarrow$

$$e = e_0 = \frac{(q_N - q_I)L_i L_j}{L_i + L_j}, \quad (23)$$

then any strategy profile $\{\alpha I + (1-\alpha)N; \beta I + (1-\beta)N; \lambda_j A_j + (1-\lambda_j)A_i\}$, with α and β set according to (22) is a mixed strategy Nash equilibrium. Recall that (15) must hold.

We can see that when $\lambda_i \neq \lambda_j$, the conditions in (17)-(18) and (20)-(21) dictate that only one user plays a mixed strategy at a time while the other plays a pure strategy. Moreover, the attacker chooses the value of λ that corresponds to the user playing the mixed strategy. This consideration is critical to understand the next two cases.

Case M2: If $\lambda_j < \lambda_i \Rightarrow$

$$e < e_0 = \frac{(q_N - q_I)L_i L_j}{L_i + L_j}, \quad (24)$$

and $\lambda = \lambda_i$, then according to (21), User j plays the pure strategy I. This means that $\beta = 1$. Setting $\beta = 1$ in (22) yields

$$\alpha = \alpha_0 = \frac{q_N(L_i + \pi L_j) - q_I(L_j + \pi L_i)}{(q_N - q_I)(L_i + \pi L_j)}. \quad (25)$$

We can verify that $0 < \alpha_0 < 1$ when $\pi > \pi_0$ and (1), (2) and (3) hold. Therefore, the strategy profile $\{\alpha_0 I + (1-\alpha_0)N; I; \lambda_i A_j + (1-\lambda_i)A_i\}$ is a mixed strategy Nash equilibrium. Observe that the low profile User i is more likely to invest in this mixed strategy Nash equilibrium compared to the pure strategy Nash equilibrium (N, I, A_j). In this scenario, it is relatively cheap to invest in security as shown in (24).

However, If $\lambda_j < \lambda_i$ and $\lambda = \lambda_j$, then according to (18) User i plays the pure strategy I. This means that $\alpha = 1$. Setting $\alpha = 1$ in (22) yields

$$\beta = \frac{q_N(L_j + \pi L_i) - q_I(L_i + \pi L_j)}{(q_N - q_I)(L_i + \pi L_j)} > 1. \quad (26)$$

The last Inequality in (26) holds when (1), (2), and (3) holds. This is a contradiction with (15).

Case M3: If $\lambda_j > \lambda_i \Rightarrow$

$$\frac{(q_N - q_I)L_i L_j}{L_i + L_j} < e < (q_N - q_I)L_i. \quad (27)$$

Note that the last inequality must hold because of (4). Thus according to (17), when $\lambda = \lambda_j$, User i plays the pure strategy N. This means that $\alpha = 0$. Setting $\alpha = 0$ in (22) yields:

$$\beta = \beta_0 = \frac{q_N[(L_j + \pi L_i) - (L_i + \pi L_j)]}{(q_N - q_I)(L_j + \pi L_i)}. \quad (28)$$

We can verify that $0 < \beta_0 < 1$ when $\pi > \pi_0$ and (1), (2) and (3) hold. Therefore, the strategy profile $\{N; \beta_0 I + (1-\beta_0)N; \lambda_j A_j + (1-\lambda_j)A_i\}$ is a mixed strategy Nash equilibrium. Observe that the high profile User j is less likely to invest in this mixed strategy Nash equilibrium compared to the pure strategy Nash equilibrium (N, I, A_j). In this scenario, it is relatively more expensive to invest in security as shown in (27).

However, If $\lambda_j > \lambda_i$ and $\lambda = \lambda_i$, then according to (20), User j plays the pure strategy N. This means that $\beta = 0$. Setting $\beta = 0$ in (22) yields:

$$\alpha = -\frac{q_N[(L_j + \pi L_i) - (L_i + \pi L_j)]}{(q_N - q_I)(L_i + \pi L_j)} < 0 \quad (29)$$

The last inequality in (29) holds when (1), (2), and (3) hold. This is a contradiction with (15). ■

In summary, we have shown that the low profile User i imposes two different types of negative externalities on the high profile User j in the cloud. If L_i is low enough in such a way that (14) holds and $\pi \leq \pi_0$, then in the pure strategy profile (N, I, A_j) shown in subcase (4a), the attacker targets the high profile user even though the high profile user (User j) invests in security while the low profile user (User i) does not invest. User j is the attacker's only target. This is the first type of negative externality. When L_i is high enough in such a way that (14) does not hold, then $\pi > \pi_0$and the attacker is forced to play a mixed strategy. The specific mixed strategy is determined by the total expense required to invest in security e. However, User i produces the second type of negative externality by investing less often than User j in all those mixed strategies. In fact, there is no Nash equilibrium in which the low profile user (User i) plays the pure strategy I.

Furthermore, with low value of e (Case M2), it can be shown that User i's probability to invest α_0 (see (25)) increases with L_i to the benefit of User j. Recall that in Case M2, User j always invests. However, if the value of e is high (Case M3), it is easy to verify that User j probability to invest in security β_0 (see (28)) decreases with L_i. Recall that in Case M3, User i does not invest (play N). A high value of e causes an under investment problem in cloud security.

In short, it is important for a high profile user to be collocated with other high profile users in a public cloud. The notion of externality has always being perceived in the housing market. In fact, the value of other homes in the same neighborhood influences the price of any particular home. As a consequence, a rational home buyer will try to find out who are his neighbors before buying a home. A similar concept should apply to cloud computing. It can be important that a cloud user knows who his neighbors are. A cloud user's neighborhood is the set of users with whom he shares the same resources (hypervisor, CPU cycle, DRAM of the physical machine, physical memory, and network buffers).

VI. NUMERICAL RESULTS

Our game analysis has provided a detailed exposition of our game model and its equilibrium properties. The numerical results in this section are derived from our game analysis. The main variables used in calculating pure and mixed strategy equilibrium were $R, q_I, q_N, L_i, L_j, \pi,$ and e. We will use specific numbers to provide concrete examples and examine the three cases in which we will increase $e, L_j,$ and π individually while *ceteris paribus*.

A. Changes in User *j*'s Payoff with Probability π

In this first scenario, we will take the value of π to be variable while setting values for all the other parameters. We will take $q_N = 0.5, q_I = 0.1, R = 1.2, L_i = 1, L_j = 10$. Those values are chosen to illustrate some of the non-intuitive implications of our game model. Using (11), we can see that $\pi_0 = 0.102$. Furthermore, with (23) we can see that $e_0 = 0.3636$. Moreover, (27) gives us $0.3636 < e < 0.4$. Recall that in case of a mixed strategy Nash equilibrium ($\pi > \pi_0 = 0.102$), the value of e determines which of the mixed strategy Nash equilibrium (Case M1, M2 or M3) is selected by the players. In Fig. 2, we set $e = 0.3$ ($e < e_0$) so that once the critical value of π is reached, the mixed strategy Nash equilibrium will be as Case M2.

We immediately see that the payoff for User j in pure Nash equilibrium is negative. When the payoff of a rational user is negative, he prefers not to use the cloud. So, for all values of $\pi \leq 0.102$ the User j, which is assumed to be rational in our model, will not use the cloud because the risk of a security breach and negative externalities of using the cloud are greater than the multiple benefits that cloud computing provides. Recall that in the pure strategy Nash equilibrium, User j is at a disadvantage because he is the attacker's only target.

However, at $\pi = 0.102$ there is a strategy change from pure to mixed due to (11), and as at this point the strategies shift. With a shift in Nash equilibrium and players' strategies, there is a concurring change in the function used as it is a new set of equations governing the strategies. This allows for a positive payoff for $0.102 < \pi \leq 0.47837$ and implies that User j will participate in the cloud for the aforementioned values of π. These results are seemingly counterintuitive since the hypervisor has a higher probability of being compromised when User j participates in cloud activities than when he does not. This is explained by the equilibrium shift to a mixed strategy where the attacker is not only attacking User j but also User i. This lowers User j's potential loss and thus shifts his payoff upwards.

Examining Fig. 2 again, the payoff becomes negative again as π crosses 0.47837, which shows that User j will again not participate in the cloud for all values of $0.47837 < \pi \leq 1$ since the probability of being compromised from an indirect attack is now too high to justify cloud usage.

By setting $e = 0.38$ and upholding (27), Fig. 3 shows the strategy shift from pure Nash equilibrium to the mixed Nash equilibrium in Case M3. Still, for values of $\pi \leq 0.102$, User j will not participate in the cloud because of his negative payoff. Although once π crosses 0.102, a change in payoff from negative to positive, as in Fig. 2, makes the cloud a viable option. Interestingly, the payoff does not cross over again to become negative after this original movement of equilibriums. This means that for all values of $0.102 < \pi \leq 1$, User j will participate in the cloud if $0.3636 < e < 0.4$. Another surprising result is that User j's payoff is higher in Fig. 3 compared to Fig. 2 although the required expense in security e in Fig. 3 is higher. Fig. 4 and 5 show more details in the change of User j's payoff with e.

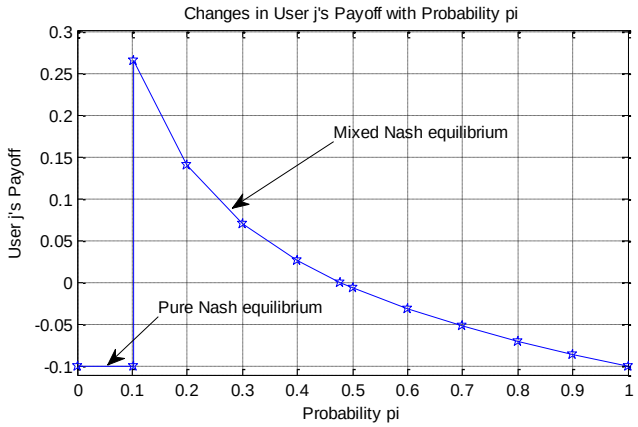

Figure 2: Changes in User j's payoff with probability π with $e < e_0$.

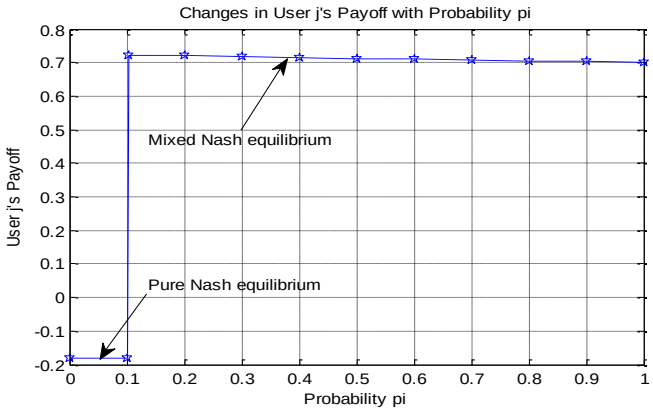

Figure 3: Changes in User j's payoff with probability π with $e > e_0$.

B. Changes of User *j*'s Payoff with the Expense in Security *e*

We have already examined the case of pure Nash equilibrium and 2 cases of mixed strategy equilibrium dependent on the varying values of π. We will now make π a constant while varying the levels of e. As stated before, the value of $\pi_0 = 0.102$ is a focal point between mixed and pure strategy equilibrium. In this case of $\pi \leq 0.102$, User j has only one (pure) strategy, whose payoff of $R - e - q_I L_j$ yields the linear function in Fig. 4.

The "x" intercept where the payoff is 0 (at $e = 0.2$) is yet another turning point where User j will no longer use the cloud. For values $0 \leq e \leq 0.2$, User j will participate in the cloud because of the low overhead of investing in security. However, for $e > 0.2$, the cost is too great to allow for a

positive payoff and User j will not use the cloud. For .102< $\pi \leq 1$ the players' strategies are switched and the entire payoff map changes as seen in Fig. 5.

In Fig. 5, we have set $\pi = 0.11 > \pi_0$ and thus we can see the three different cases of mixed strategy: Case M2 ($e <$ 0.3636), Case M1 ($e = 0.3636$) and Case M3 ($3636 < e <$ .4). The major shift from Case M2 to Case M3 occurs at the threshold of $e = 0.3636$ (Case M1) due to (23) stated in the previous analysis. For $0 \leq e < 0.3636$, the change from using to not using the cloud occurs at $e = 0.08606$ when the payoff becomes negative.

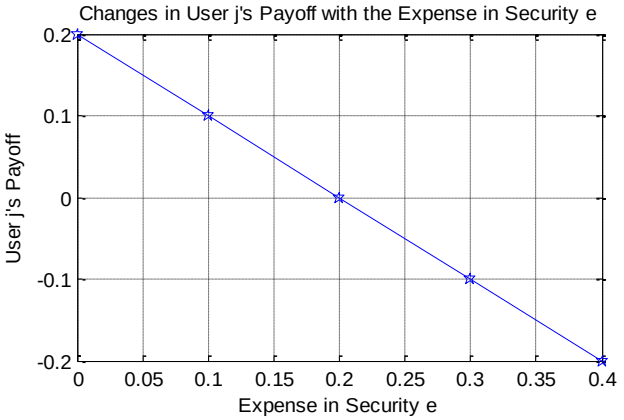

Figure 4: Changes of User j's payoff with the expense in security e with $\pi < \pi_0$.

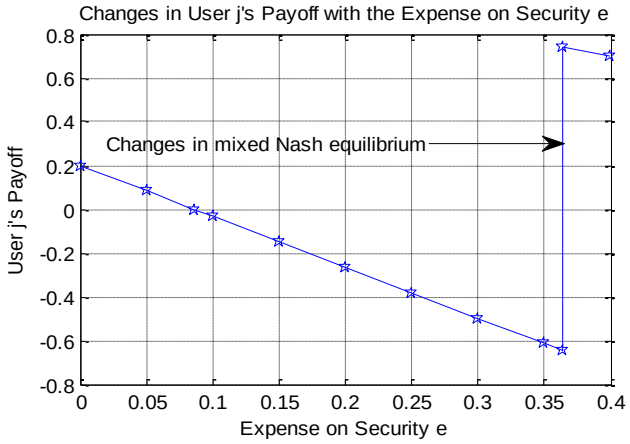

Figure 5: Changes of User j's payoff with the expense in security e with $\pi > \pi_0$.

When the expense e increases and $0.3636 < e < 0.4$, the shift in mixed Nash equilibrium from Case M2 to Case M3 causes the payoff to change and become positive. Thus it becomes possible for User j to profitably use cloud services. This is a counter intuitive result from this analysis. One may expect an increase of the expense e to never benefit User j. However, in this game theoretic setting, User j's payoff depends not only of his own action but also on the action of User i and the attacker. The increase of the expense e changes User i's and the attacker's strategy is in such a way that it has an overall positive effect on User j's payoff. In Case M3, User j invests with probability β_0 as opposed to 1 in Case M2. This yields some savings that increases User j's overall payoff. Recall that moving from Case M2 to M3 changes the mixed strategy Nash equilibrium from $\{\alpha_0 I + (1-\alpha_0)N; I; \lambda_i A_j + (1-\lambda_i)A_i\}$ to $\{N; \beta_0 I + (1-\beta_0)N; \lambda_j A_j + (1-\lambda_j)A_i\}$. Note also that for $e \geq 0.4$, Case M3 no longer applies as consistent with (4).

C. *Changes in User j's payoff with his loss from security breach L_j*

Now that the variability of π and e - and their resulting equilibrium shifts they cause - have been examined, we will examine Fig. 6 at the phenomena in equilibrium changes associated with varying values of L_j. Since L_j is a variable in both the equations that govern the values of π_0 (Equation (11)) and e_0 (Equation (23)), we must set specific values for π and e in order to avoid a problem of double variables. For the rest of the analysis of L_j, we will set $\pi = 0.1$ and $e = 0.3$. Recall that we have set $L_i = 1$. Therefore, L_j is a direct indication of how much time L_j is bigger than L_i.

Unlike the previous two problems in which a certain change in the discrete value of π with a varying e could cause an equilibrium shift, there is no such change here. Here the values of π and e are constant and L_j is the unique variable. As can be seen in Fig. 6, any value of $L_j \geq 9.8$ will result in a pure Nash equilibrium due to (11). Further, (23) shows that when $3 < L_j < 9.8$ the mixed strategy Nash equilibrium profile of Case M2 will hold, Case M1 holds for $L_j = 3$, and if $1 < L_j < 3$, then Case M3 will be used.

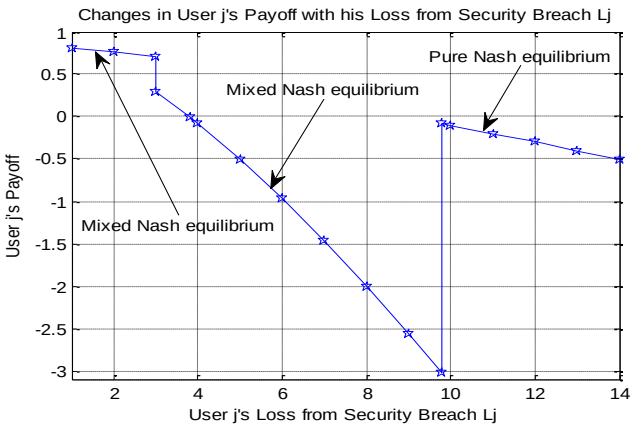

Figure 6: Changes in User j's payoff with his loss from security breach L_j.

These results show that Case M3 is the "best" of all the equilibriums because User j's potential loss L_j is so close to User i's loss L_i. An obvious result is that User j's payoff is maximized in Case M3 when L_j is close to $L_i = 1$. That is because there is no imbalance between L_i and L_j and thus the negative externalities are minimized. The negative externality in a public cloud security can be mitigated by putting VMs that have similar potential loss from a security breach in the same physical machine. However, a surprising

result is that User *j*' s payoff jumps up concurrent with switching from the mixed Nash equilibrium (Case M2) to the pure Nash equilibrium despite the fact that L_j becomes substantially greater than L_i. For instance, User *j*'s payoff when $L_j = 4L_i$ equals User *j*'s payoff when $L_j = 10L_i$. This prediction is not possible without a thorough game theoretic analysis.

D. Changes in User j's payoff with his reward from using the cloud

For the constant R, changing it will have a trivial effect on any of the given graphs shown. As seen in Fig. 7, a change in the value of R will cause the graph to translate upward or downward depending on the new value of R selected. For this particular instance, if the reward for using the cloud is increased from 1.2 to 4.4, the entire payoff scheme from $1 \leq L_j \leq 14$ becomes positive since the increased level of reward increases the payoff.

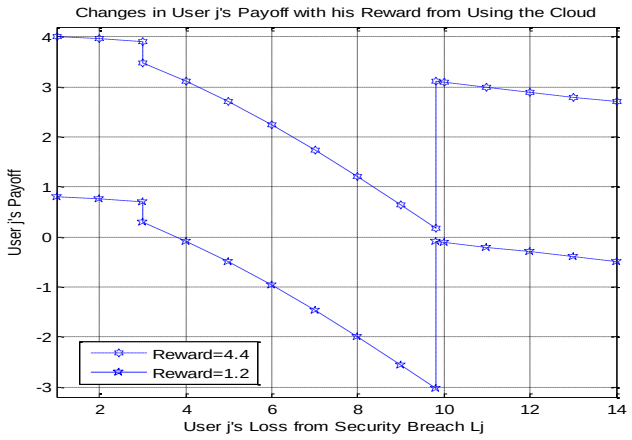

Figure 7: Changes in User *j*'s payoff with his reward from using the cloud.

VII. Model Extension and Discussion

The model we have presented so far has considered two users and one attacker. However, our model can be extended to more than two users and multiple attackers.

A. Model Extension to more than two Users and a Single Attacker

All the assumption made in our game model in Section IV remains valid except that we increase the number of users from *2* to *n*. The *n* users are denoted User *1*, User *2*, …, User *n-1*, User *n*. Their potential loss from a security breach is L_1, L_2, …,L_{n-1}, L_n respectively. We consider that $L_1 \leq L_2 \leq \cdots \leq L_{n-1} \leq L_n$. The attacker targets one of the *n* users. A similar analysis as above shows that the game admits a pure strategy Nash equilibrium if L_n is substantially greater than L_{n-1}. In this Nash equilibrium, User *n* is the attacker's only target. The attacker plays the strategy A_n, User *n* invests (plays *I*) while all the other users do not invest (play *N*). Regarding the threshold value of π below for which we have a pure strategy Nash equilibrium, (11) translates to

$$\pi_0^* = \frac{q_I L_n - q_N L_{n-1}}{q_N L_n - q_I L_{n-1}}. \qquad (30)$$

As before, the game admits a multitude of mixed strategies if $\pi > \pi_0^*$. The expense *e* will determine the specific mixed strategy the players choose.

B. Model Extension to more than two Users and Multiple Attacker

In a game with multiple independent attackers, each attacker maximizes his own payoff. If $\pi < \pi_0^*$, each attacker plays the strategy A_n and User *n* invests (plays *I*) while all the other users do not invest (play *N*). However, the game complexity increases if the attackers collude by coordinating their action and sharing the payoff. Nevertheless, an increase in the number of attackers increases the likelihood that a given user can be targeted by one attacker and eventually get compromised. As the number of attackers increases, the cloud environment becomes more hostile and more users will be forced to invest (because of (4) and (5)).

Another consideration is the users' payoff structure. There are applications in which a user incurs the same loss after being compromised by a single attacker or multiple attackers *e.g.*, information integrity can be lost when either a few bits or when many bits of a data item become useless. Either critical data are well protected or they are not. However, the severity of other types of attacks such as a Distributed Denial of Service (DDoS) increases with the number of attackers involved.

VIII. Conclusion

The lack of an accurate evaluation of the negative externalities stemming from a high profile organization using the cloud could result in the refusal of such organizations from joining a public cloud in spite of the many advantages that cloud computing offers. The negative externalities of using a public cloud come from the fact that the users are not perfectly isolated from one another. They share common resources such as the hypervisor, the last-level cache (LLC), memory bandwidth, and IO buffers that cause interdependency.

This research has used game theory to provide a quantitative approach to perform a cost benefit analysis of cloud services while taking into account the action of other cloud users and their different potential losses from a security breach. Our model takes into account the potential collateral damage from an indirect attack and cross side channel attack. The game has multiple possible Nash equilibria that can be in pure or mixed strategy. Our research finds that an increase in the probability that the hypervisor is compromised, given a successful attack on a user's VM, may force the small cloud participant to protect their VM and thus increases the overall cloud security to yield better outcome to high profile users.

This research has also shown that there is an intricate relationhip between the total expenses required to invest in security and a high profile user's payoff. A change in security expense changes the game Nash equilibria that the

players adopt with some of those equilibria being more desirable to high profile users.

Definitely, the negative externality in a public cloud security can be mitigated by putting VMs that have similar potential loss from a security breach in the same physical machine.

According to Ross Anderson, information security is hard because defenders have to defend everywhere and attackers could attack anywhere [16]. This leads to many problems for: network defenders, users, for software used in critical infrastructure, a small business, or a division in the United States government. Moreover, these security problems are exacerbated when using cloud computing. By utilizing game theory, we can more accurately describe the nature of the attacker and his motives. However, sometimes our best friend can be our worst enemy. Other players' behaviors can be seemingly erratic and even counterintuitive, which can be very dangerous when your decisions are based on the decisions of others. With game theory, we can quell some of this contradictory behavior that is characteristic of network security and bring clarity to this complex topic.

ACKNOWLEDGMENT

This research was performed while Dr. Joon Park and Dr. Manuel Rodriguez held a National Research Council (NRC) Research Associateship Award and Dr. Ming Zhao held a Summer Faculty Award at the Air Force Research Laboratory (AFRL). This research was supported by the Air Force Office of Scientific Research (AFOSR) and the Department of Defense grant 62705-CS-REP.

REFERENCES

[1] Handbook, Handbook, Occupational Outlook, Bureau of labor statistics, United States Department of Labor, Spring (2008).

[2] D. Clemente, "*Cyber Security and Global Interdependence: What is Critical?"*, Chatham House, 2013.

[3] K. Cukier, *"Ensuring and Insuring Critical Information Infrastructure Protection: A Report of the 2005 Rueschlikon Conference on Information Policy,"* The Rueschlikon Conference, 2005.

[4] F. Hare, *"The Interdependent Nature of National Cyber Security: Motivating Private Action for a Public Good,"* PhD Dissertation, School of Public Policy, George Mason University, (2011).

[5] R. Myerson (1991). "*Game Theory: Analysis of Conflict,"* Harvard University Press, p. 1.

[6] G. Heal, H. Kunreuther. *"You only die once: Managing discrete interdependent risks,"* No. w9885. National Bureau of Economic Research, 2003.

[7] H. Kunreuther, H. Geoffrey *"Interdependent Security: the Case of Identical Agents,"* Working paper, Columbia Business School and Wharton Risk Management and Decision Processes Center. Journal of Risk and Uncertainty, forthcoming, Special Issue on Terrorist Risks, 2002.

[8] W. Sun, X. Kong, D. He, X. You. *"Information security problem research based on game theory,"* International Symposium on Publication Electronic Commerce and Security, 2008.

[9] C. Kamhoua, N. Pissinou, K. Makki. *"Game theoretic modeling and evolution of trust in autonomous multi-hop networks: Application to network security and privacy,"* IEEE International Conference on Communications (ICC), 2011.

[10] T. Alpcan, T. Başar. *"Network security: A decision and game-theoretic approach,"* Cambridge University Press, 2010.

[11] N. Leavitt, *"Is cloud computing really ready for prime time,"* Growth 27.5 (2009).

[12] P. Mell, T. Grance. *"The NIST definition of cloud computing (draft),"* NIST special publication 800.145 (2011): 7.

[13] S. Pearson, A. Benameur. *"Privacy, security and trust issues arising from cloud computing,"* IEEE Second International Conference on Cloud Computing Technology and Science (CloudCom) 2010.

[14] Zissis, Dimitrios, and Dimitrios Lekkas. *"Addressing cloud computing security issues,"* Future Generation Computer Systems 28.3 (2012): 583-592.

[15] C. Everett, *"Cloud computing–A question of trust,"* Computer Fraud & Security 2009.6 (2009): 5-7.

[16] R. Anderson, *"Why Information Security is Hard – an Economic Perspective,"* Working paper, Computer Laboratory, Cambridge. 2001

http://www.ftp.cl.cam.ac.uk/ftp/users/rja14/econ.pdf

[17] J. Horrigan, *"Use of cloud computing applications and services,"* Pew Internet & American Life project memo, September 2008.

[18] T. Ristenpart, E. Tromer, H. Shacham, S. Savage. *"Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds,"* In the proceedings of the 16th ACM Conference on Computer and Communications Security, CCS'09, Chicago, IL, USA, October 2009.

[19] A. Bates, B. Mood, J. Pletcher, H. Pruse, M. Valafar, K. Butler *"Detecting Co-Residency with Active Traffic Analysis Techniques,"* in the proceedings of the 2012 ACM Cloud Computing Security Workshop (CCSW) in conjunction with the 19th ACM Conference on Computer and Communications Security, October 2012, Raleigh, North Carolina, USA.

[20] Y. Zhang, A. Juels, A. Oprea, M. Reiter *"HomeAlone: Co-Residency Detection in the Cloud via Side-Channel Analysis,"* in the proceedings of IEEE Symposium on Security and Privacy, May 2011, Oakland, California, USA.

[21] L. Carin, G. Cybenko, J. Hughes, *"Cybersecurity Strategies: The QuERIES Methodology,"* Computer, vol.41, no.8, pp.20-26, Aug. 2008.

[22] United States Securities and Exchange Commission

http://www.sec.gov/divisions/corpfin/guidance/cfguidance-topic2.htm

[23] C. Li, A. Raghunathan, N. Jha, *"A Trusted Virtual Machine in an Untrusted Management Environment,"* IEEE Transactions on Services Computing, vol. 5, no. 4, pp. 472-483, Fourth Quarter 2012.

[24] A. Azab, P. Ning, Z. Wang, X. Jiang, X. Zhang, N. Skalsky *"HyperSentry: enabling stealthy in-context measurement of hypervisor integrity,"* In Proceedings of the 17th ACM conference on Computer and communications security (CCS '10). ACM, New York, NY, USA.

[25] Federal Register / Vol. 78, No. 33 / Tuesday, February 19, 2013 / Presidential Documents

[26] P. Tailor, L. Jonker *"Evolutionary Stable Strategies and Game Dynamic,"* Mathematical Biosciences, 5:455-484, 1978.

[27] Charles Kamhoua, Luke Kwiat, Kevin Kwiat, Joon Park, Ming Zhao, Manuel Rodriguez, "Game Theoretic Modeling of Security and Interdependency in a Public Cloud" in the proceedings of IEEE International Conference on Cloud Computing, (IEEE CLOUD 2014) Anchorage, Alaska, June 2014.

Charles A. Kamhoua (S'10–M'12–SM'14) received his B.S. in Electronic from the University of Douala (ENSET), Cameroon in 1999, and the M.S. in Telecommunication and Networking and PhD in Electrical Engineering from Florida International University in 2008 and 2011 respectively. In 2011, he joined

the Cyber Assurance Branch of the U.S. Air Force Research Laboratory, Rome, New York, as a National Academies Postdoctoral Fellow, and in 2012 became a Research Electronics Engineer. His current research interests cover the application of game theory and mechanism design to cyber security and survivability, with over 30 technical publications. He is a reviewer of multiple journals and serves on the technical program committees of several international conferences. Dr. Kamhoua has won numerous prestigious awards, including an Air Force Notable Achievement Award, a Best Paper Award at the 2013 International Symposium on Foundations of Open Source Intelligence and Security Informatics (FOSINT-SI 2013), a National Science Foundation (NSF) PIRE award at Fluminense Federal University, Brazil, the Air Force Office of Scientific Research (AFOSR) Windows on the World Visiting Research Fellowship at Oxford University, UK and an AFOSR basic research award. He is an advisor for the National Research Council, a member of the National Society of Black Engineer (NSBE) and a Senior Member of IEEE.

Luke A. Kwiat is currently an honors student at the University of Florida in Gainesville where he is pursuing his Bachelors of Science in Industrial Engineering. He has held engineering internships, through the Griffiss Institute, at the Air Force Research Laboratory in Rome, New York. His research interests are in applying economic and game-theoretic methods to the optimization of engineering endeavors. His work has been published in the Proceedings of IEEE International Conference on Cloud Computing (IEEE CLOUD 2014). He is a member of the Sigma Nu fraternity.

Kevin A. Kwiat is a Principal Computer Engineer with the U.S. Air Force Research Laboratory (AFRL) in Rome, New York where he has worked for over 31 years. Currently is assigned to the Cyber Assurance Branch. He received the BS in Computer Science and the BA in Mathematics from Utica College of Syracuse University, and the MS in Computer Engineering and the Ph.D. in Computer Engineering from Syracuse University. He holds 4 patents. In addition to his duties with the Air Force, he is an adjunct professor of Computer Science at the State University of New York at Utica/Rome, an adjunct instructor of Computer Engineering at Syracuse University, and a Research Associate Professor with the University at Buffalo. He is an advisor for the National Research Council. He has been by recognized by the AFRL Information Directorate with awards for best paper, excellence in technology teaming, and for outstanding individual basic research. His main research interest is dependable computer design.

Dr. Joon S. Park is an associate professor at the School of Information Studies (iSchool), Syracuse University, Syracuse, New York, USA. Over the past decades Prof. Park has been involved with theoretical/practical research, education, and services in information and systems security. He served as the founding director of the Certificate of Advanced Study (CAS) in Information Security Management (ISM) at the Syracuse iSchool (2003-2013). He is Syracuse University's Point of Contact (POC) at the Center of Academic Excellence (CAE) in Information Assurance / Cyber Defense programs, which are designated by the National Security Agency (NSA) and the Department of Homeland Security (DHS).

Ming Zhao is an associate professor of the School of Computing and Information Sciences (SCIS) at Florida International University (FIU), where he directs the research laboratory for Virtualized Infrastructures, Systems, and Applications (VISA). His research interests include virtualization, cloud computing, high-performance systems, and autonomic computing. His research has been funded by the National Science Foundation (NSF), Department of Homeland Security, Department of Defense, and industry companies, and his research outcomes have been adopted by several production systems in the industry. He has received the NSF Faculty Early Career Development (CAREER) award, the Air Force Summer Faculty Fellowship, the Air Force Visiting Faculty Fellowship, the VMware Faculty Award, the FIU SCIS Excellence in Student Mentoring award, and the Best Paper award of the IEEE International Conference on Autonomic Computing. He received his B.E. and M.E. in Automation from Tsinghua University, and his Ph.D. in Electrical and Computer Engineering from the University of Florida. He is a member of IEEE and the IEEE Computer Society.

Dr. Manuel Rodriguez is a senior researcher at the Air Force Research Laboratory. From 1997 to 2002, he was a member of the Dependable Computing and Fault Tolerance (TSF) group in the research laboratory LAAS-CNRS (Toulouse, France). He received the M.S. degree (1998) from the Polytechnic University of Valencia (UPV), Spain, and the Ph.D. degree (2002) from the National Polytechnic Institute of Toulouse (INPT), France. Dr. Rodriguez has participated in projects sponsored by government (AFOSR, NASA, ESA), and industry. His research interests include security, dependability & fault tolerance, testing, formal methods, real-time systems, and software & hardware reliability.